Remco den Heijer (4626184)

Go Deep or Go Home?

Responsible Professors: David Tax, Marco Loog
Supervisors: Tom Viering, Ziqi Wang, Taylan Turan, Yuko Kato

1. Introduction

Do convolutional neural networks (CNNs) have to be deep to learn regression tasks?

Relevance

- Deep networks are generally harder to train [1]
- Contribute to understanding of key blocks in deep regression [2]

Regression tasks

- Image -> mean pixel value
- Image -> median pixel value
- Image -> standard deviation of pixel values

2. Evaluate performance on tasks

Experiment

- Fine-tune learning rate
- Train and test on FashionMNIST
- Repeat 10 times

Hyperparameters

- Batch size:64
- Epochs: 30

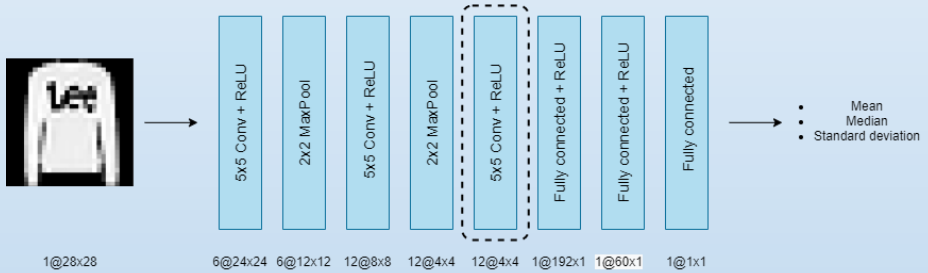

3. Average test losses per task and model

	MeanFMNIST	MedianFMNIST	StdFMNIST
Loss (MSE)	0.016	0.071	0.0061

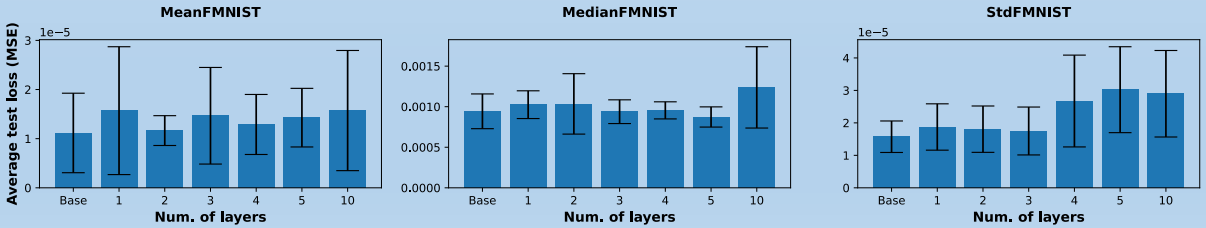

4. Discussion and conclusion

Conclusions

- Deep CNNs are not required to learn these regression tasks
- Going deeper does not improve performance

Discussion

- Early stopping of training
- Simple regression tasks that have no real world value

[1]: Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Training very deep networks. *arXiv preprint arXiv:1507.06228, 2015.*
[2]: Stéphane Lathuilière, Pablo Mesejo, Xavier Alameda-Pineda, and Radu Horaud. A comprehensive analysis of deep regression. *IEEE transactions on pattern analysis and machine intelligence*, 42(9):2065–2081, 2019.